

Formal Concept Analysis of Rodent Carriers of Zoonotic Disease

Roman Ilin

Sensors Directorate, Bldg 620., Wright Patterson AFB, OH 45433 USA

RILIN325@GMAIL.COM

Barbara A. Han

Cary Institute of Ecosystem Studies, Box AB Millbrook, NY. USA

HANB@CARYINSTITUTE.ORG

Abstract

The technique of Formal Concept Analysis is applied to a dataset describing the traits of rodents, with the goal of identifying zoonotic disease carriers, or those species carrying infections that can spillover to cause human disease. The concepts identified among these species together provide rules-of-thumb about the intrinsic biological features of rodents that carry zoonotic diseases, and offer utility for better targeting field surveillance efforts in the search for novel disease carriers in the wild.

1. Introduction

Formal Concept Analysis (FCA) is a data mining technique built on the foundation of lattice theory (Ganter, 1999). FCA is applied to data tables with binary features (contexts) resulting in a hierarchical representation of the data (concept lattice) consisting of frequent patterns of feature co-occurrence, called concepts. Due to recent computing and algorithmic advances, it is now possible to compute concept lattices for real world data sets (Andrews, 2009; 2015). Here, FCA was applied to contexts of wild rodents to identify concepts indicative of zoonotic disease carriers, or those species carrying infections that can spillover to cause human disease. The concepts identified among these species together provide rules-of-thumb about the intrinsic biological features of rodents that carry zoonotic diseases, and offer utility for better targeting field surveillance efforts in the search for novel disease carriers in the wild.

This work builds on the analysis presented in (Han et al., 2015) where a machine learning technique was applied in order to build a predictive model of carrier species. The dataset consisted of representatives of positive (carrier) and

negative (non-carrier) species. The generalized boosted regression analysis built a classifier and simultaneously identified the top predictive features. The overall classification accuracy was 90

Although this method was able to identify important predictive features across all species, further analysis was necessary to understand the interactions between these features, and to identify motifs of shared features that are common among subsets of positive species. An approach to such analysis is given in this contribution by utilizing FCA. FCA was conducted on binarized positive and negative contexts. We were able to identify particular biological concepts shared among rodents that carry zoonotic diseases. FCA also identified particular features for which additional empirical data collection would disproportionately improve the capacity to predict novel disease reservoirs, illustrating the kind of discourse between data mining and empirical data collection that will benefit infectious disease surveillance.

The rest of this paper is organized as follows. Section 2 provides background for FCA. Section 3 describes the rodent data and Random Forest results. Section 4 describes the application of FCA. We discuss the results and outline future plans in Section 5.

2. Formal Concept Analysis

Formal Concept Analysis (FCA) is a data mining technique built on the foundation of lattice theory. This section provides a brief mathematical description. See (Davey, 2002; Ganter, 1999) for more details.

Suppose that U is a finite set of objects, and V is a finite set of properties (also called attributes or features). Suppose that a binary relation $R \subseteq U \times V$ is defined. The expression xRy means that object x has feature y . The triplet $K = (U, V, R)$ is called a formal context.

We can express the set of features for an object x , as the set $xR \equiv y \in V | xRy$. Similarly, the set of objects with some

property y is the set $Ry \equiv x \in U | xRy$.

The following two mappings between the sets U and V are defined:

$$A^* \equiv \bigcap_{x \in A} xR, A \subseteq U \quad (1)$$

$$B^* \equiv \bigcap_{y \in B} Ry, B \subseteq V \quad (2)$$

A^* is the set of features shared by all objects in the set A . Similarly, B^* is the set of objects that all share the same set of features. An ordered pair (A, B) , $A \subseteq U, B \subseteq V$, is called a **formal concept**, if $A^* = B$ and $B^* = A$. The set of objects in the formal concept is referred to as the **extension**, and the set of features as the **intension** of the concept. For a concept C , the extension and the intension are denoted as $Ext(C)$ and $Int(C)$. It is a classical result in FCA that the mappings between U and V form Galois connection.

The mappings

$$X^{**} : 2^U \rightarrow 2^U \quad (3)$$

$$Y^{**} : 2^V \rightarrow 2^V \quad (4)$$

are closure operators. A fixpoint of the closure operator, which is a set such that $X^{**} = X$, is called a closed element. As evident from the definition of formal concept, the set of intensions of all concepts is equivalent to the set of all elements closed with respect to the mapping (3), and similarly, the set of all extensions is the set of closed sets with respect to (4). We denote the set of all concepts of a given context as

$$B \equiv \{(A, B) | A^* = B \wedge B^* = A\} \quad (5)$$

and the sets of all intensions and extensions as $INT(B)$ and $EXT(B)$.

Consider the set B of all concepts of a given context. We define a relation of partial order of this set, as follows:

$$(A_1, B_1) \preceq (A_2, B_2) \equiv A_1 \subseteq A_2 \equiv B_1 \supseteq B_2 \quad (6)$$

This relation defines a partially ordered set that is a complete lattice (Davey, 2002), referred to as **concept lattice**. Suppose that T is a set of indices. The meet and join of concepts from the set $(A_t, B_t), t \in T$, are defined as follows:

$$\bigwedge_{t \in T} (A_t, B_t) \equiv \left(\left(\bigcap_{t \in T} A_t \right), \left(\bigcup_{t \in T} B_t \right)^{**} \right) \quad (7)$$

$$\bigvee_{t \in T} (A_t, B_t) \equiv \left(\left(\bigcup_{t \in T} A_t \right)^{**}, \left(\bigcap_{t \in T} B_t \right) \right) \quad (8)$$

Note that the intersection of concept extensions is an extension of some concept, but the union of extensions is not necessarily an extension of some concept. The same is true for intensions: the intersection of intensions is an intension of some concept, but not the union. We will denote the transformation of a context into the concept lattice as

$$B \equiv CL(K) \quad (9)$$

An illustration of FCA is shown in Fig. 1 for a simple context. The diagram was computed using a software package called Galicia (Gal, 2016). The lattice diagram is drawn by placing the more generic concepts (those containing fewer features) above more specific concepts (with more features).

Each concept is characterized by a number between 0

$$Support(C) \equiv \frac{|Ext(C)|}{|U|} \times 100\% \quad (10)$$

Fig. 1 shows the support for each concept in a simple context. The greater the support, the more prominent a given concept. Also, it is easy to show that support increases as we move up in the lattice, with support of the top concept being 100%. This property led to the development of an Iceberg lattice (Stumme, 2002). For a given concept lattice, the Iceberg lattice contains only the concepts with support above a certain threshold. This is equivalent to cutting off the bottom of the lattice. The remaining elements contain more generic concepts that explain a large portion of the context. In the example in Fig. 1 the 80% Iceberg would contain the top two concepts (disregarding the trivial top concept). We observe that 80% of all the objects contain one or both of these concepts.

3. Rodent Data

Rodent trait data were obtained from PanTHERIA, a species-level database of life history, ecological, and geographical traits of the world's mammals (Jones, 2009). The traits include features such as adult body mass, age of first birth, litter size; geographical region given by the maximum and the minimum latitude and longitude, etc. The total number of predictive features was 88. Values are




Figure 1. Illustration of Formal Concept Analysis. Data table contains 5 objects and 4 attributes. The concept lattice consists of 6 concepts. A concept can be thought of as a square in the data table, which is illustrated for one of the concepts.

the result of numerous field studies, thus there are many missing values in the dataset.

The rodent dataset contained 2277 objects (species), with 217 labelled as positive (disease reservoirs), and the rest as negative (reservoir status unknown). The analysis in (Han et al., 2015) utilized boosted regression trees to train classifiers, using up to 10,000 trees and 10-fold cross-validation to prevent overfitting. Analysis was done using gbm package in R (Ridgeway, 2006). This analysis identified top 15 predictor features (Table 1).

4. Analysis of Concepts in Rodent Data

Formal Concept Analysis was conducted on the reduced dataset containing 2277 objects and 15 features. Since the features are integer or real valued, they had to be discretized. For non-geographical features, we computed the median values, and mapped each value to one of three categories: high (above median), low (below median), and

Table 1. Top 15 predictor features for Rodent Data

X26.1_GR_Area_km2	GEOGRAPHIC AREA, SQUARE KM.
X23.1_SexualMaturityAge_d	AGE OF SEXUAL MATURITY, DAYS
X27.2_HuPopDen_Mean_n.km2	HUMAN POPULATION DENSITY, km^{-2}
logNeoBM	NEONATAL BODY MASS, LOG(GRAMS)
X15.1_LitterSize	LITTER SIZE
X5.1_AdultBodyMass_g	ADULT BODY MASS, G
X9.1_GestationLen_d	GESTATION LENGTH, DAYS
X25.1_WeaningAge_d	WEANING AGE, DAYS
X13.1_AdultHeadBodyLen_mm	ADULT LENGTH HEAD AND BODY, MM
SpeciesDensity	DENSITY OF MAMMAL SPECIES, km^{-2}
X30.2_PET_Mean_mm	MEAN POTENTIAL EVAPOTRANSPIRATION RATE, MM
X26.2_GR_MaxLat_dd	MAXIMUM LATITUDE OF THE GEOGRAPHICAL RANGE
X16.1_LittersPerYear	NUMBER OF LITTERS PER YEAR
X26.5_GR_MaxLong_dd	MAXIMUM LONGITUDE OF THE GEOGRAPHICAL RANGE
X26.3_GR_MinLat_dd	MINIMUM LATITUDE OF THE GEOGRAPHICAL RANGE

NAN (missing). For geographical features, latitude was separated into 4 bins: NAN (missing), [-90, -30], [-30, 30], and [30, 90]. Longitude was separated into 3 bins: NAN (missing), [-180, -25], and [25, 180]. Biogeographical boundaries roughly correspond to the tropical vs. subtrop-

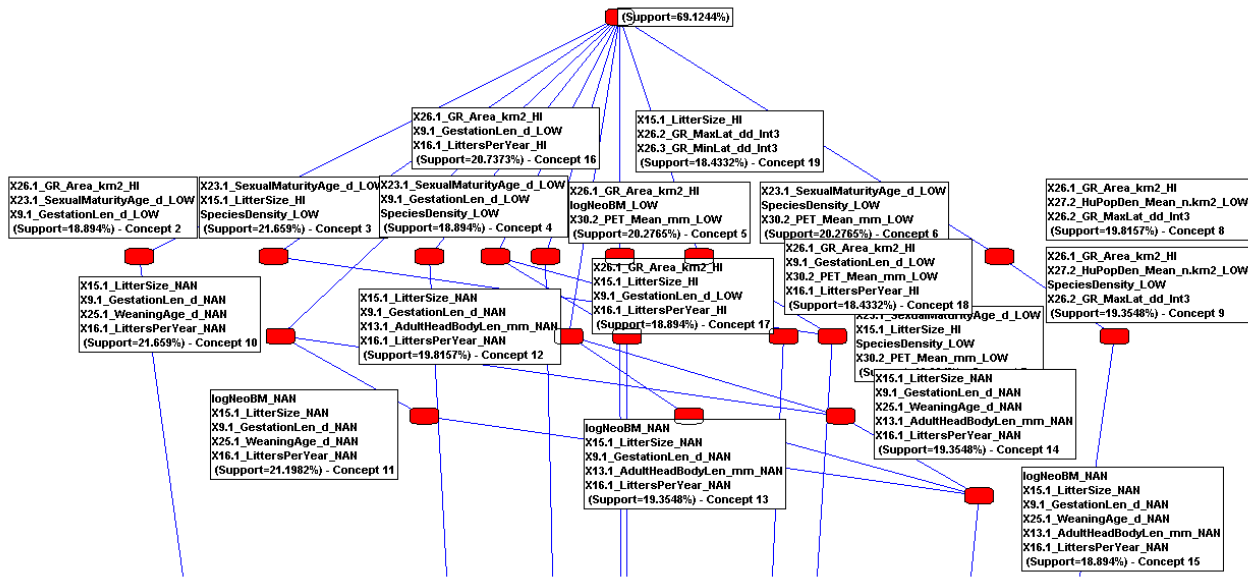


Figure 2. Iceberg Lattice for Rodent Context at 18% minimum support.

ical regions and Eurasia vs. the Americas. The resulting binary dataset contained 47 columns, one for each feature category.

Formal Concept Analysis was conducted using In-Close software (Andrews, 2009). The positive context, with 217 objects, resulted in 6,197 concepts. The negative concepts, with 2,060 objects, resulted in 137,515 objects. Large number of concepts is typical for FCA. Iceberg analysis was applied to the large number of concepts, which is typical of FCA, to focus on more general concepts present in the data and identify strong dependencies between features.

Each concept is a pattern of feature co-occurrence in the data. If a certain pattern is found in both the negative and the positive lattice, it alone cannot be used to decide whether a certain object belongs to a positive or negative class. Based on this reasoning, we removed all concepts from the positive lattice that also exist in the negative class. The reduced positive lattice with 738 concepts was used for further analysis. We computed the Iceberg lattice with minimum support of 18%. The diagram is shown in Fig. 2. This lattice has 18 concepts, and 6 of them are the missing data concepts those with names ending in NAN. These concepts are supported by 69% of the objects in the positive class.

5. Conclusion

This contribution outlines the initial results of applying FCA to the rodent data. We provided a brief description

of FCA as a mathematical framework for frequent pattern search in categorical datasets. FCA improves existing methods for identifying traits characterizing potential zoonotic disease carriers by exploring the dependencies between different traits. FCA also identifies particular data items that are disproportionately represented among disease carriers that should be prioritized for future field work. For example, we found that about 21% of all positive species contain a pattern of large litter size, early age at sexual maturity, and living in areas with high mammal biodiversity (species density). This pattern is not encountered in the negative species, suggesting that more species with this pattern need to be tested for zoonotic diseases. We also found that about 21% of positive species contained a pattern of missing data on weaning age, litter size, length of gestation period, and the number of litters per year. However, these data were available for negative species. This contrast highlights a particular need to measure this suite of features over others in future fieldwork to make the greatest improvements in the prediction of positive species. In future work, we plan to expand the analysis by utilizing more concepts from FCA. For example, we plan to evaluate stability scores of concepts (Kuznetsov et al., 2007). We also plan to expand our analysis to other datasets, in particular to the analysis of bats to identify concepts describing species that carry filoviruses, which cause hemorrhagic fevers such as Ebola virus disease.

References

- Formal concept analysis, university of montreal,
<http://www.iro.umontreal.ca/galicia/>, 2016.
- Andrews, Simon. In-close, a fast algorithm for computing formal concepts. 2009.
- Andrews, Simon. A 'best-of-breed' approach for designing a fast algorithm for computing fixpoints of galois connections. *Information Sciences*, 2015.
- Davey, B.; Priestley, H. *Introduction to Lattices and Order*. Cambridge: University Press, 2002.
- Ganter, B.; Wille, R. *Formal Concept Analysis: Mathematical Foundations*. New York: Springer-Verlag, 1999.
- Han, B A, Schmidt, J P, Bowden, S E, and Drake, J M. Rodent reservoirs of future zoonotic diseases. *PNAS*, 2015.
- Jones, K. E. Pantheria: A species-level database of life history, ecology, and geography of extant and recently extinct mammals. *Ecology*, 2009.
- Kuznetsov, Sergei, Obiedkov, Sergei, and Roth, Camille. Reducing the representation complexity of lattice-based taxonomies. In *Conceptual Structures: Knowledge Architectures for Smart Applications*, pp. 241–254. Springer, 2007.
- Ridgeway, G. Generalized boosted regression models. documentation on the r package gbm, 2006.